


AF\$
IFW



IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Application No.: 10/616,437
Filed: July 9, 2003
Inventor(s):
Hans F. van Rietschote, Craig W.
Hobbs, and Mahesh P. Saptarshi

Title: Migrating Virtual Machines
among Computer Systems
to Balance Load Caused by
Virtual Machines

§ Examiner: Truong, Camquy
§ Group/Art Unit: 2195
§ Atty. Dkt. No: 5760-12200
§
§
§
§
§
§
§
§
§
§

I hereby certify that this correspondence is being deposited with the United States Postal Service with sufficient postage as first class mail in an envelope addressed to Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450, on the date indicated below.

Lawrence J. Merkel
Printed Name
[Signature] 3/2/06
Signature Date

APPEAL BRIEF

Mail Stop Appeal Brief - Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir/Madam:

Further to the Notice of Appeal filed January 9, 2006 (with Pre-Appeal Brief Request for Review) and the Pre-Appeal Brief Decision mailed February 21, 2006, Appellants present this Appeal Brief. Appellants respectfully request that this appeal be considered by the Board of Patent Appeals and Interferences.

I. REAL PARTY IN INTEREST

The present application is owned by Symantec Corporation, through merger with Veritas Software Corporation and Veritas Operating Corporation. An assignment of the present application to Veritas Operating Corporation is recorded at Reel 014291, Frame 0283.

II. RELATED APPEALS AND INTERFERENCES

There are no related appeals or interferences known to Appellant.

III. STATUS OF CLAIMS

Claims 1-33 are pending. Claims 1-33 are rejected under 35 U.S.C. § 103(a) and the rejection of these claims under section 103(a) is being appealed. A copy of claims 1-33 is included in the Claims Appendix attached hereto.

IV. STATUS OF AMENDMENTS

No unentered amendment to the claims has been filed after final rejection.

V. SUMMARY OF CLAIMED SUBJECT MATTER

Independent claim 1 is directed to a cluster comprising a plurality of computer systems (10A-10N). Each of the plurality of computer systems is configured to execute one or more virtual machines (16A-16E), and each computer system comprises hardware (20A-20N) and a plurality of instructions (38A). When the instructions are executed on the hardware, they detect that a first load of a first computer system of the plurality of computer systems exceeds a second load of a second computer system of the plurality of computer systems. The instructions migrate at least a first virtual machine executing on the first computer system to the second computer system responsive to detecting that the first load exceeds the second load. (See, e.g., Figs. 2 and 3; specification, page 8, line 22-

page 9, line 16; and specification page 13, line 7-page 17, line 28).

Independent claim 14 is directed to a method. One or more virtual machines (16A-16E) are scheduled for execution on hardware (20A) comprising a first computer system (10A) of a plurality of computer systems (10A-10N). The first computer system detects that the first computer system has a first load that exceeds a second load of a second computer system of the plurality of computer systems. Responsive to the detecting, the first computer system migrates at least a first virtual machine executing on the first computer system to a second computer system (10B) of the plurality of computer systems. (See, e.g., Figs. 2 and 3; specification, page 8, line 22-page 9, line 16; and specification page 13, line 7-page 17, line 28).

Independent claim 23 is directed to a computer accessible medium (120) encoded with a plurality of instructions (38A). The instructions, when executed on a first computer system (10A) select a first virtual machine (16A) from one or more virtual machines (16A-16E) to be scheduled for execution on the first computer system responsive to a first load of the first computer system exceeding a second load of a second computer system (10B) of a plurality of computer systems (10A-10B) including the first computer system. The instructions, when executed, migrate the first virtual machine to the second computer system to be executed on the second computer system. (See, e.g., Figs. 2, 3, and 8; specification, page 8, line 22-page 9, line 16; specification page 13, line 7-page 17, line 28; and specification, page 23, lines 2-20).

VI. GROUND OF REJECTION TO BE REVIEWED ON APPEAL

1. Claims 1-11 and 13-33 are rejected under 35 U.S.C. § 103(a) as being unpatentable over alleged Applicant Admitted Prior Art ("AAPA") in view of Saito et al., U.S. Patent No. 6,578,064 ("Saito").
2. Claim 12 is rejected under 35 U.S.C. § 103(a) as being unpatentable over AAPA in view of Saito and Bodin et al., U.S. Patent No. 5,675,762 ("Bodin").

VII. ARGUMENT

First Ground of Rejection:

Claims 1-11 and 13-33 are rejected under 35 U.S.C. § 103(a) as being unpatentable over AAPA in view of Saito. Appellants traverse this rejection for at least the following reasons.

Appellants note that the rejections at issue in this appeal all rely on alleged AAPA. Appellants do not agree that the AAPA is prior art. The passages from the background of the present application that are relied on in the rejections do not rise to a clear, unequivocal admission of prior art. The rejection is also in error as highlighted below.

Claims 1-2, 13, 23-25, and 33:

Appellants respectfully submit that each of claims 1, 14, and 23 recites a combination of features not taught or suggested in AAPA and Saito. For example claim 1 recites a combination of features including: "detects that a first load of a first computer system of the plurality of computer systems exceeds a second load of a second computer system of the plurality of computer systems and migrates at least a first virtual machine executing on the first computer system to the second computer system responsive to detecting that the first load exceeds the second load".

The Office Action alleges that the AAPA teaches "migrates at least a first virtual machine executing on the first computer system to the second computer system" at page 1, lines 17-21. However, that passage states "Virtual machines may also be used to provide high availability. If a given computer system fails, tasks that were executing in virtual machines on that computer system may be transferred ("failed over") to similar virtual machines on other computer systems. Moreover, the entire virtual machine may be failed over to another computer system." Accordingly, the AAPA describes transferring virtual machines away from a computer system **when that computer system fails**. Since the system has failed, **it is not possible that a failover is "responsive to**

detecting that the first load exceeds the second load". In fact, since the computer system described in the AAPA has failed, there is really no load at all on that computer system. The failover is responsive to the failure of the computer system, not the load on the computer system.

The teachings of Saito cannot be use to overcome this deficiency in the teachings of the AAPA. To do so would change the principle of operation of the AAPA from a high availability system, that recovers from the failure of a computer system, to a load balancing system. Nothing in Saito teaches or suggests use in a high availability system. In essence, by teaching the detection of failure in a computer system and failing over to another system responsive to the failure, the AAPA teaches away from the alleged combination since no load can be measured when a failure has occurred. Furthermore, load balancing according to the teachings of Saito may even interfere with the high availability of operation of the AAPA. For example, if the load balancing operation does not leave enough available resources on one system to fail over another system, the AAPA high availability operation is defeated.

For at least the above stated reasons, Appellants submit that the rejection of claim 1 is unsubstantiated in the art, and the rejection should be withdrawn. Claims 2 and 13 depend from claim 1, and thus the rejection of these claims is not substantiated in the art and should be withdrawn as well for at least the above stated reasons. Each of claims 2 and 13 recites additional combinations of features not taught or suggested in the cited art.

Claim 23 recites a combination of features including: "select a first virtual machine ... responsive to a first load of the first computer system exceeding a second load of a second computer system of a plurality of computer systems including the first computer system; and migrate the first virtual machine to the second computer system to be executed on the second computer system". The same teachings of the AAPA and Saito highlighted above with regard to claim 1 are alleged to teach the above highlighted features of claim 23. Appellants respectfully submit that the AAPA and Saito do not teach or suggest the above highlighted features, either. For at least the above stated

reasons, Appellants submit that the rejection of claim 23 is unsubstantiated in the art, and the rejection should be withdrawn. Claims 24-25 and 33 depend from claim 23, and thus the rejection of these claims is not substantiated in the art and should be withdrawn as well for at least the above stated reasons. Each of claims 24-25 and 33 recites additional combinations of features not taught or suggested in the cited art.

For at least all of the above stated reasons, Appellants respectfully submit that the rejection of claims 1-2, 13, 23-25, and 33 is in error and request reversal of the rejection.

Claims 3 and 26:

Claims 3 and 26 depend from claims 1 and 23, respectively. Accordingly, Appellants respectfully submit that the rejection of these claims is in error for at least the reasons stated above with regard to claims 1 and 23. Additionally, claim 3 recites a combination of features including: "the plurality of instructions, when executed on the first computer system, select the second computer system to compare loads".

The Office Action alleges that the above features are taught in Saito, citing col. 4, lines 18-22 and col. 24, lines 38-49. Appellants respectfully disagree. Saito does generally teach determining loads among the computer systems. However, Saito is not specific as to how such determination is performed. It is clear, however, that Saito compares loads among all computers: "Program migrate determination procedure means 296, 297 determine, through mutual communications between the computers, whether loads on the respective computers are equally distributed in accordance with their performances. If a load on a computer containing the program migrate determination procedure means is larger as compared with loads on other computers, the program migrate determination procedure means determines to migrate a program in the associated computer so as to bring about the equal load distribution." (Saito, col. 24, lines 25-33). The method of communication is unspecified, and thus there is no teaching or suggestion for the first computer system to select the second computer system to compare loads. For example, the computers could broadcast their loads to all other computers, and then the computers could compare the received loads from all other computers to their

own load.

Claim 26 recites a combination of features including: "the plurality of instructions, when executed, select the second computer system to compare loads." The same teachings of Saito highlighted above with regard to claim 3 are alleged to teach the features of claim 26. Appellants respectfully submit that Saito does not teach or suggest the features of claim 26, either.

For at least all of the above stated reasons, Appellants respectfully submit that the rejection of claims 3 and 26 is in error and request reversal of the rejection.

Claims 4 and 27:

Claims 4 and 27 depend from claims 3 and 26, respectively, and thus the rejection of claims 4 and 27 is in error for at least the reasons given above for claims 3 and 26. Additionally, claim 4 recites a combination of features including: "the plurality of instructions, when executed on the first computer system, randomly select the second computer system from the plurality of computer systems".

The Office Action does not separately treat the features of claim 4 from the features of claim 3. However, Appellants respectfully submit that nothing in Saito teaches or suggests randomly selecting a computer system with which to compare loads. In fact, since Saito clearly teaches that load comparison is performed between all computers, there would be no need to randomly select computers for comparison. Rather, random selection would be counter-productive to producing a comparison of all loads among all computers.

Claim 27 recites a combination of features including: "wherein the second computer system is randomly selected from the plurality of computer systems". The same teachings of Saito highlighted above are alleged to teach the above features of claim 27. Appellants respectfully submit that Saito does not teach or suggest the features of claim 27, either.

For at least the above stated reasons, Appellants respectfully submit that the rejection of claims 4 and 27 is in error and request reversal of the rejection.

Claims 5-6 and 28-29:

Claims 5-6 and 28-29 depend from claims 1 and 23, respectively, and thus the rejection of claims 1 and 23 is in error for at least the reasons given above for claims 1 and 23. Additionally, claim 5 recites a combination of features including: "the first virtual machine has a corresponding load that is nearest, among loads of the virtual machines executing on the first computer system, to 1/2 the difference between the first load and the second load".

The Office Action alleges that Saito teaches the above highlighted features at col. 22, lines 26-47. Appellants respectfully disagree. Saito teaches: "Since n is now set to one, the flow returns 25 to step 322, where the load/performance values are calculated for the respective computers, and the difference value D is next calculated at step 323. Although the load/performance values for the computers designated 102 and subsequent reference numerals are omitted in FIG. 13, assuming that the load/performance value for the computer 100 equal to 1.10 is the largest and the load/performance value for the computer 101 equal to 0.67 is the smallest when n is one, the difference D is calculated to be 0.43. Therefore, the difference D exceeds the threshold Dt when n is one. It is determined that the load is not equally distributed when the difference D exceeds the threshold Dt, so that a program to be migrated is determined (step 325)." (Saito, col. 22, lines 26-39). Thus, Saito teaches calculating the difference between loads (D) and determining that a program is to migrate if that raw difference exceeds a threshold (Dt). This does not teach or suggest "the first virtual machine has a corresponding load that is nearest, among loads of the virtual machines executing on the first computer system, to 1/2 the difference between the first load and the second load" as recited in claim 5.

Claim 28 recites a combination of features including: "the first virtual machine has a corresponding load that is approximately 1/2 the difference between the first load

and the second load." The same teachings of Saito highlighted above with regard to claim 5 are alleged to teach the above features of claim 28. Appellants respectfully submit that Saito does not teach or suggest the above highlighted features of claim 28, either.

For at least the above stated reasons, Appellants respectfully submit that the rejection of claims 5 and 28 is in error and request reversal of the rejection. Claims 6 and 29, being dependent from claims 5 and 28, respectively, recite additional combinations of features not taught or suggested in AAPA and Saito, and thus the rejection of these claims is in error and reversal is respectively requested.

Claims 7 and 30:

Claims 7 and 30 depend from claims 5 and 28, respectively, and thus the rejection of claims 7 and 30 is in error for at least the reasons given above for claims 5 and 28. Additionally, each of claims 7 and 30 recite a combination of features including: "the corresponding load of the first virtual machine represents a target load programmed for the first virtual machine in the first computer system"

The Office Action alleges that Saito teaches the above highlighted features, citing col. 24, lines 28-49 and col. 25, lines 3-4. However, col. 24, lines 28-49 discusses the computers trading information on loads in general, and selecting a program to migrate. Col. 25, lines 3-4 discuss improving real time performance of the programs. This has nothing to do with a target load programmed for the first virtual machine in the first computer system.

For at least the above stated reasons, Appellants respectfully submit that the rejection of claims 7 and 30 is in error and request reversal of the rejection.

Claims 8 and 31:

Claims 8 and 31 depend from claims 1 and 23, respectively, and thus the rejection of claims 8 and 31 is in error for at least the reasons given above for claims 1 and 23.

Additionally, each of claims 8 and 31 recite a combination of features including: "the first virtual machine has a first corresponding load on the first computer system and a second corresponding load on the second computer system, and wherein the first corresponding load differs from the second corresponding load".

The Office Action alleges that the AAPA teaches the above highlighted features, citing page 1, lines 23-29. However, these teachings are: "Typically, virtual machines are manually assigned to computer systems by an administrator, based on an assumed load that each virtual machine will generate or based on measured loads under various circumstances. However, estimates of the load may be incorrect or the load may dynamically change during use, and thus less than optimal usage of the computer systems may result. In some cases, cluster-level performance monitoring tools may be used to monitor load and inform the administrator when load imbalances are detected. Complicated policy settings may be used to help balance the load." Nothing in this section teaches or suggests that the first corresponding load (of the first virtual machine on the first computer system) differs from the second corresponding load (of the first virtual machine on the second computer system).

For at least the above stated reasons, Appellants respectfully submit that the rejection of claims 8 and 31 is in error and request reversal of the rejection.

Claims 9-10 and 32:

Claims 9-10 and 32 depend from claims 1 and 23, respectively, and thus the rejection of claims 9-10 and 32 is in error for at least the reasons given above for claims 1 and 23. Additionally, claim 9 recites a combination of features including: "the first virtual machine has a corresponding load that is calculated as a weighted combination of measurements of usage of two or more resources of the first computer system".

The Office Action alleges that the above highlighted features are taught in Saito at col. 25, lines 37-41. Appellants respectfully disagree. Saito teaches "the load amount may also be represented by the sum of processing amounts or predicted processing time

of executable programs. The latter way is advantageous in that loads on respective computers are more accurately evaluated so that fully equal load distribution can be accomplished." (Saito, col. 25, lines 39-44) Nothing in this section teaches or suggests weighting the various amounts that Saito would sum.

For at least the above stated reasons, Appellants respectfully submit that the rejection of claims 9 and 32 is in error and request reversal of the rejection. Claim 10, dependent from claim 9, recites an additional combination of features not taught or suggested in the cited art. Accordingly, the rejection of claim 10 is also in error and Appellants respectfully request reversal.

Claim 11:

Claim 11 depends from claim 9, and thus the rejection of claim 11 is in error for at least the reasons given above for claim 9. Additionally, claim 11 recites a combination of features including: "the measurements of usage include an amount of input/output activity generated by the first virtual machine during execution".

The Office Action alleges that Saito teaches the above features at col. 25, lines 55-65. However, these teachings are: "Furthermore, while the distributed computing system according to the third or fourth embodiment evaluates an amount of load on a computer based on the number of executable programs with the urgency level ranging from zero to an urgency level given to a program to be initiated or migrated, this evaluation may be made based on the number of executable programs with the same urgency level as that of a program to be initiated or migrated. This evaluation is advantageous in that the time required to calculate the load amount is reduced so that a more efficient operation can be accomplished." Nothing in this section has anything to do with an amount of I/O activity generated by the first virtual machine during execution.

For at least the above stated reasons, Appellants respectfully submit that the rejection of claim 11 is in error and request reversal of the rejection.

Claims 14-15 and 22:

Claim 14 recites a combination of features including: "the first computer system detecting that the first computer system has a first load that exceeds a second load of a second computer system of the plurality of computer systems; and the first computer system migrating at least a first virtual machine executing on the first computer system to a second computer system of the plurality of computer systems responsive to the detecting". The same teachings of the AAPA and Saito highlighted above with regard to claim 1 are alleged to teach the above highlighted features of claim 14. Appellants respectfully submit that the AAPA and Saito do not teach or suggest the above highlighted features, either.

Furthermore, claim 14 recites "the first computer system detecting that the first computer system has a first load that exceeds a second load of a second computer system...and the first computer system migrating at least a first virtual machine ... responsive to the detecting". Clearly, the first computer system in AAPA cannot perform the above since, under the AAPA, the first computer system has failed to trigger the fail over.

For at least the above stated reasons, Appellants submit that the rejection of claim 14 is unsubstantiated in the art, and the rejection should be withdrawn. Claims 15 and 22 depend from claim 14, and thus the rejection of these claims is not substantiated in the art and should be withdrawn as well for at least the above stated reasons. Each of claims 15 and 22 recites additional combinations of features not taught or suggested in the cited art.

For at least all of the above reasons, Appellants respectfully submit that the rejection of claims 14-15 and 22 is in error and request reversal of the rejection.

Claim 16:

Claim 16 depends from claim 14, and thus the rejection of claim 16 is in error for at least the reasons given above with regard to claim 14. Additionally, claim 16 recites a

combination of features including: "selecting the second computer system to compare loads."

The Office Action alleges that claim 16 is rejected for the same reasons as claims 2 and 14. However, these teachings have nothing to do with selecting a computer system to compare loads. Additionally, as highlighted above with regard to claims 3 and 26, Saito does generally teach determining loads among the computer systems. However, Saito is not specific as to how such determination is performed. It is clear, however, that Saito compares loads among all computers: "Program migrate determination procedure means 296, 297 determine, through mutual communications between the computers, whether loads on the respective computers are equally distributed in accordance with their performances. If a load on a computer containing the program migrate determination procedure means is larger as compared with loads on other computers, the program migrate determination procedure means determines to migrate a program in the associated computer so as to bring about the equal load distribution." (Saito, col. 24, lines 25-33). The method of communication is unspecified, and thus there is no teaching or suggestion to select the second computer system to compare loads. For example, the computers could broadcast their loads to all other computers, and then the computers could compare the received loads from all other computers to their own load.

For at least all of the above reasons, Appellants respectfully submit that the rejection of claim 16 is in error and request reversal of the rejection.

Claim 17:

Claim 17 depends from claim 16, and thus the rejection of claim 17 is in error for at least the reasons given above for claim 16. Additionally, claim 17 recites a combination of features including: "selecting the second computer system is random."

Appellants respectfully submit that nothing in Saito teaches or suggests randomly selecting a computer system with which to compare loads. In fact, since Saito clearly teaches that load comparison is performed between all computers (as highlighted above

with regard to claim 16), there would be no need to randomly select computers for comparison. Rather, random selection would be counter-productive to producing a comparison of all loads among all computers.

For at least all of the above reasons, Appellants respectfully submit that the rejection of claim 17 is in error and request reversal of the rejection.

Claim 18:

Claim 18 depends from claim 17, and thus the rejection of claim 18 is in error for at least the reasons given above for claim 17. Additionally, claim 18 recites a combination of features including: "each of the plurality of computer systems periodically randomly selecting another one of the plurality of computer systems to compare loads and to potentially migrate virtual machines."

As highlighted above with regard to claim 17, Saito does not teach or suggest randomly selecting the second computer. For similar reasons, Saito does not teach or suggest periodically randomly selecting a computer. Additionally, the Office Action alleges that Saito teaches the features of claim 18 at col. 23, lines 25-37. Appellants respectfully disagree. This section of Saito refers to repeating the comparing of loads and migration of a program over all urgency levels (see, e.g., Saito col. 23, lines 33-37). This has nothing to do with the above highlighted features of claim 18.

For at least all of the above reasons, Appellants respectfully submit that the rejection of claim 18 is in error and request reversal of the rejection.

Claim 19:

Claim 19 depends from claim 14, and thus the rejection of claim 19 is in error for at least the reasons given above for claim 14. Additionally, claim 19 recites a combination of features including: "the first virtual machine has a corresponding load that is nearest, among the virtual machines executing on the first computer system, to 1/2 the difference between the first load and the second load."

The Office Action alleges that Saito teaches the above highlighted features at col. 22, lines 26-47. Appellants respectfully disagree. Saito teaches: "Since n is now set to one, the flow returns 25 to step 322, where the load/performance values are calculated for the respective computers, and the difference value D is next calculated at step 323. Although the load/performance values for the computers designated 102 and subsequent reference numerals are omitted in FIG. 13, assuming that the load/performance value for the computer 100 equal to 1.10 is the largest and the load/performance value for the computer 101 equal to 0.67 is the smallest when n is one, the difference D is calculated to be 0.43. Therefore, the difference D exceeds the threshold Dt when n is one. It is determined that the load is not equally distributed when the difference D exceeds the threshold Dt, so that a program to be migrated is determined (step 325)." (Saito, col. 22, lines 26-39). Thus, Saito teaches calculating the difference between loads (D) and determining that a program is to migrate if that raw difference exceeds a threshold (Dt). This does not teach or suggest "the first virtual machine has a corresponding load that is nearest, among loads of the virtual machines executing on the first computer system, to 1/2 the difference between the first load and the second load" as recited in claim 19.

For at least all of the above reasons, Appellants respectfully submit that the rejection of claim 19 is in error and request reversal of the rejection.

Claim 20:

Claim 20 depends from claim 14, and thus the rejection of claim 20 is in error for at least the reasons given above for claim 14. Additionally, claim 20 recites a combination of features including: "the first virtual machine has a first corresponding load on the first computer system and a second corresponding load on the second computer system, and wherein the first corresponding load differs from the second corresponding load".

The Office Action alleges that the AAPA teaches the above highlighted features, citing page 1, lines 23-29. However, these teachings are: "Typically, virtual machines

are manually assigned to computer systems by an administrator, based on an assumed load that each virtual machine will generate or based on measured loads under various circumstances. However, estimates of the load may be incorrect or the load may dynamically change during use, and thus less than optimal usage of the computer systems may result. In some cases, cluster-level performance monitoring tools may be used to monitor load and inform the administrator when load imbalances are detected. Complicated policy settings may be used to help balance the load." Nothing in this section teaches or suggests that the first corresponding load (of the first virtual machine on the first computer system) differs from the second corresponding load (of the first virtual machine on the second computer system).

For at least all of the above reasons, Appellants respectfully submit that the rejection of claim 20 is in error and request reversal of the rejection.

Claim 21:

Claim 21 depends from claim 14, and thus the rejection of claim 21 is in error for at least the reasons given above for claim 14. Additionally, claim 21 recites a combination of features including: "calculating the load of the first virtual machine as a weighted combination of measurements of usage of two or more resources of the first computer system."

The Office Action alleges that the above highlighted features are taught in Saito at col. 25, lines 37-41. Appellants respectfully disagree. Saito teaches "the load amount may also be represented by the sum of processing amounts or predicted processing time of executable programs. The latter way is advantageous in that loads on respective computers are more accurately evaluated so that fully equal load distribution can be accomplished." (Saito, col. 25, lines 39-44) Nothing in this section teaches or suggests weighting the various amounts that Saito would sum.

For at least all of the above reasons, Appellants respectfully submit that the rejection of claim 21 is in error and request reversal of the rejection.

Second Ground of Rejection:

Claim 12 is rejected under 35 U.S.C. § 103(a) as being unpatentable over the alleged AAPA in view of Saito and Bodin. Appellants traverse this rejection for at least the following reasons.

Claim 12:

Claim 12 depends from claim 9, and thus the rejection of claim 12 is in error for at least the reasons given above for claim 9. Additionally, claim 12 recites a combination of features including: "the measurements of usage include an amount of memory occupied by the first virtual machine."

The Office Action alleges that the above highlighted features are taught in Bodin in the abstract and at col. 2, lines 41-45. Appellants respectfully disagree. Bodin does teach selectively updating the page directory to track memory usage in the memory block used by the virtual machine (col. 2, lines 43-45). However, the usage being tracked is the usage by the peripheral device (col. 2, lines 40-43): "Memory may also be locked down in order to permit the peripheral device to access that section of memory." Thus, the cited section refers to peripheral device memory usage and does not teach or suggest the above highlighted features of claim 12.

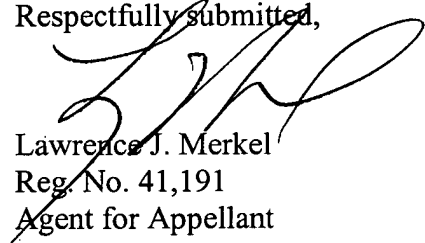
For at least all of the above reasons, Appellants respectfully submit that the rejection of claim 12 is in error and request reversal of the rejection.

VIII. CONCLUSION

For the foregoing reasons, it is submitted that the Examiner's rejection of claims 1-33 under 35 U.S.C. § 103(a) was erroneous, and reversal of the decision is respectfully requested.

The Commissioner is authorized to charge the appeal brief fee of \$500.00 and any other fees that may be due to Meyertons, Hood, Kivlin, Kowert, & Goetzel, P.C. Deposit Account No. 501505/5760-12200/LJM. This Appeal Brief is submitted with a return receipt postcard.

Respectfully submitted,



Lawrence J. Merkel
Reg. No. 41,191
Agent for Appellant

Meyertons, Hood, Kivlin, Kowert & Goetzel, P.C.
P.O. Box 398
Austin, TX 78767-0398
(512) 853-8850

Date: March 2, 2006

IX. CLAIMS APPENDIX

The claims on appeal are as follows.

1. A cluster comprising a plurality of computer systems, wherein each of the plurality of computer systems is configured to execute one or more virtual machines, each of the plurality of computer systems comprising hardware and a plurality of instructions that, when executed on the hardware, detects that a first load of a first computer system of the plurality of computer systems exceeds a second load of a second computer system of the plurality of computer systems and migrates at least a first virtual machine executing on the first computer system to the second computer system responsive to detecting that the first load exceeds the second load.
2. The cluster as recited in claim 1 wherein the first virtual machine executes on the second computer system independent of the first computer system, even if the first virtual machine was initially launched on the first computer system.
3. The cluster as recited in claim 1 wherein the plurality of instructions, when executed on the first computer system, select the second computer system to compare loads.
4. The cluster as recited in claim 3 wherein the plurality of instructions, when executed on the first computer system, randomly select the second computer system from the plurality of computer systems.
5. The cluster as recited in claim 1 wherein the first virtual machine has a corresponding load that is nearest, among loads of the virtual machines executing on the first computer system, to $1/2$ the difference between the first load and the second load.
6. The cluster as recited in claim 5 wherein the corresponding load of the first virtual machine represents the actual load experienced in executing the first virtual machine on the first computer system.

7. The cluster as recited in claim 5 wherein the corresponding load of the first virtual machine represents a target load programmed for the first virtual machine in the first computer system.
8. The cluster as recited in claim 1 wherein the first virtual machine has a first corresponding load on the first computer system and a second corresponding load on the second computer system, and wherein the first corresponding load differs from the second corresponding load, wherein the first computer system is configured to transmit one or more load factors to the second computer system, and wherein the second computer system is configured to calculate the second corresponding load from the one or more load factors, and wherein the first computer system and the second computer system are configured to exchange the first corresponding load and the second corresponding load to select the first virtual machine for migration.
9. The cluster as recited in claim 1 wherein the first virtual machine has a corresponding load that is calculated as a weighted combination of measurements of usage of two or more resources of the first computer system.
10. The cluster as recited in claim 9 wherein the measurements of usage include an amount of time that the first virtual machine is executing in a central processing unit of the first computer system.
11. The cluster as recited in claim 9 wherein the measurements of usage include an amount of input/output activity generated by the first virtual machine during execution.
12. The cluster as recited in claim 9 wherein the measurements of usage include an amount of memory occupied by the first virtual machine.
13. The cluster as recited in claim 1 wherein each of the plurality of computer systems include a schedule having a plurality of entries, each entry corresponding to a virtual

machine to be executed on the respective one of the plurality of computer systems, and wherein migrating the first virtual machine comprises deleting the entry corresponding to the first virtual machine in the schedule of the first computer system and inserting the entry corresponding to the first virtual machine in the schedule of the second computer system.

14. A method comprising:

scheduling one or more virtual machines for execution on hardware comprising a first computer system of a plurality of computer systems;

the first computer system detecting that the first computer system has a first load that exceeds a second load of a second computer system of the plurality of computer systems; and

the first computer system migrating at least a first virtual machine executing on the first computer system to a second computer system of the plurality of computer systems responsive to the detecting.

15. The method as recited in claim 14 further comprising executing the first virtual machine on the second computer system independent of the first computer system, even if the first virtual machine was initially launched on the first computer system.

16. The method as recited in claim 14 further comprising:

selecting the second computer system to compare loads; and

selecting the first virtual machine to migrate to the second computer system responsive to the first load exceeding the second load.

17. The method as recited in claim 16 wherein selecting the second computer system is

random.

18. The method as recited in claim 17 further comprising each of the plurality of computer systems periodically randomly selecting another one of the plurality of computer systems to compare loads and to potentially migrate virtual machines.

19. The method as recited in claim 14 wherein the first virtual machine has a corresponding load that is nearest, among the virtual machines executing on the first computer system, to $1/2$ the difference between the first load and the second load.

20. The method as recited in claim 14 wherein the first virtual machine has a first corresponding load on the first computer system and a second corresponding load on the second computer system, and wherein the first corresponding load differs from the second corresponding load, the method further comprising:

the first computer system transmitting one or more load factors to the second computer system;

the second computer system calculating the second corresponding load from the one or more load factors; and

the first computer system and the second computer system exchanging the first corresponding load and the second corresponding load to select the first virtual machine for migration.

21. The method as recited in claim 14 further comprising calculating the load of the first virtual machine as a weighted combination of measurements of usage of two or more resources of the first computer system.

22. The method as recited in claim 14 wherein each of the plurality of computer systems include a schedule having a plurality of entries, each entry corresponding to a virtual

machine to be executed on the respective one of the plurality of computer systems, and wherein migrating the first virtual machine comprises:

deleting the entry corresponding to the first virtual machine in the schedule of the first computer system; and

inserting the entry corresponding to the first virtual machine in the schedule of the second computer system.

23. A computer accessible medium encoded with a plurality of instructions that, when executed on a first computer system:

select a first virtual machine from one or more virtual machines to be scheduled for execution on the first computer system responsive to a first load of the first computer system exceeding a second load of a second computer system of a plurality of computer systems including the first computer system; and

migrate the first virtual machine to the second computer system to be executed on the second computer system.

24. The computer accessible medium as recited in claim 23 wherein the plurality of instructions, when executed, schedule the one or more virtual machines for execution on hardware comprising the first computer system.

25. The computer accessible medium as recited in claim 22 wherein the first virtual machine executes on the second computer system independent of the first computer system during use, even if the first virtual machine was initially launched on the first computer system.

26. The computer accessible medium as recited in claim 23 wherein the plurality of

instructions, when executed, select the second computer system to compare loads.

27. The computer accessible medium as recited in claim 26 wherein the second computer system is randomly selected from the plurality of computer systems.

28. The computer accessible medium as recited in claim 23 wherein the first virtual machine has a corresponding load that is approximately $1/2$ the difference between the first load and the second load.

29. The computer accessible medium as recited in claim 28 wherein the corresponding load of the first virtual machine represents the actual load experienced in executing the first virtual machine on the first computer system.

30. The computer accessible medium as recited in claim 28 wherein the corresponding load of the first virtual machine represents a target load programmed for the first virtual machine in the first computer system.

31. The computer accessible medium as recited in claim 23 wherein the first virtual machine has a first corresponding load on the first computer system and a second corresponding load on the second computer system, and wherein the first corresponding load differs from the second corresponding load, and wherein the plurality of instructions, when executed:

transmit one or more load factors to the second computer system, wherein the second computer system is configured to calculate the second corresponding load from the one or more load factors; and

exchanges the first corresponding load and the second corresponding load with the second computer system to select the first virtual machine for migration.

32. The computer accessible medium as recited in claim 23 wherein the plurality of instructions, when executed, calculate the load of the first virtual machine as a weighted combination of measurements of usage of two or more resources of the first computer system.

33. The computer accessible medium as recited in claim 23 wherein each of the plurality of computer systems include a schedule having a plurality of entries, each entry corresponding to a virtual machine to be executed on the respective one of the plurality of computer systems, and wherein the plurality of instructions migrate the first virtual machine by:

deleting the entry corresponding to the first virtual machine in the schedule of the first computer system; and

inserting the entry corresponding to the first virtual machine in the schedule of the second computer system.

X. EVIDENCE APPENDIX

No evidence submitted under 37 CFR §§ 1.130, 1.131 or 1.132 or otherwise entered by the Examiner is relied upon in this appeal.

XI. RELATED PROCEEDINGS APPENDIX

There are no related proceedings known to Appellant.